

Task-related edge density (TED) - a new method for revealing large-scale network formation in fMRI data of the human brain

Gabriele Lohmann^{1,2*}, Johannes Stelzer^{1,2}, Verena Zuber³, Tilo Buschmann⁴, Daniel Margulies⁵, Andreas Bartels⁶, Klaus Scheffler^{1,2}

The formation of transient networks in response to external stimuli or as a reflection of internal cognitive processes is a hallmark of human brain function. However, its identification in fMRI data of the human brain is notoriously difficult. Here we propose a new method of fMRI data analysis that tackles this problem by considering large-scale, task-related synchronisation networks. Networks consist of nodes and edges connecting them, where nodes correspond to voxels in fMRI data, and the weight of an edge is determined via task-related changes in dynamic synchronisation between their respective time series. Based on these definitions, we developed a new data analysis algorithm that identifies edges in a brain network that differentially respond in unison to a task onset and that occur in dense packs with similar characteristics. Hence, we call this approach “Task-related Edge Density” (TED). TED proved to be a very strong marker for dynamic network formation that easily lends itself to statistical analysis using large scale statistical inference. A major advantage of TED compared to other methods is that it does not depend on any specific hemodynamic response model, and it also does not require a presegmentation of the data for dimensionality reduction as it can handle large networks consisting of tens of thousands of voxels. We applied TED to fMRI data of a fingertapping task provided by the Human Connectome Project. TED revealed network-based involvement of a large number of brain areas that evaded detection using traditional GLM-based analysis. We show that our proposed method provides an entirely new window into the immense complexity of human brain function.

¹University Hospital, Department of Biomedical Magnetic Resonance Imaging, Tübingen, Germany

²Magnetic Resonance Centre, Max-Planck-Institute for Biological Cybernetics, Tübingen, Germany

³Institute of Clinical Medicine, University of Oslo, Norway

⁴Fraunhofer Institute for Cell Therapy and Immunology, Leipzig, Germany

⁵Max-Planck-Institute for Human Cognitive and Brain Sciences, Leipzig, Germany

⁶CIN Vision & Cognition Group, Centre for Integrative Neuroscience, Tübingen, Germany

*Corresponding author: lohmann@tuebingen.mpg.de

1. Introduction

The human brain is a large-scale network consisting of approximately 85 billion neurons that form a vast number of subnetworks on all spatial scales [1]. The properties of the intrinsic connectivity, such as small-worldness [2], make possible the coexistence between local processing of information in specialised circuits and large-scale integrative processes, involving multiple remote sites. It has been suggested that the neuroanatomical architecture itself gives rise to a rich dynamic repertoire setting the frame for a large number of flexibly accessible brain functions [3, 4, 5].

Traditionally, brain mapping techniques using functional magnetic resonance imaging (fMRI) have focused on studying brain areas separately in a voxel-by-voxel fashion (i.e. univariately). The key idea behind such approaches was to identify task- or stimulus-related changes of the blood-oxygen-level dependent (BOLD) signal activity on the local level.

The most prominent example is statistical parametric mapping using the general linear model [6, 7]. Such mass-univariate approaches treat voxels *independently* from each other [8], however, they do not capture processes relating to integration and functional interplay between remote brain regions [9, 10, 11].

As a result, the focus of the neuroimaging community has shifted away from the pure segregationist to a more integrative (i.e. network-based) view [12, 13, 14, 15, 16, 17, 18, 10]. In the following, we give a brief overview over existing methods that go beyond traditional GLM-based activation maps.

Seed-based approaches investigate how the statistical dependency between a seed voxel or area changes with respect to the rest of the brain. The most prominent examples are correlation-based approaches [19], where the correlation between the time series of the seed area to all other voxels is computed. A widely used method is the *psycho-physiological*

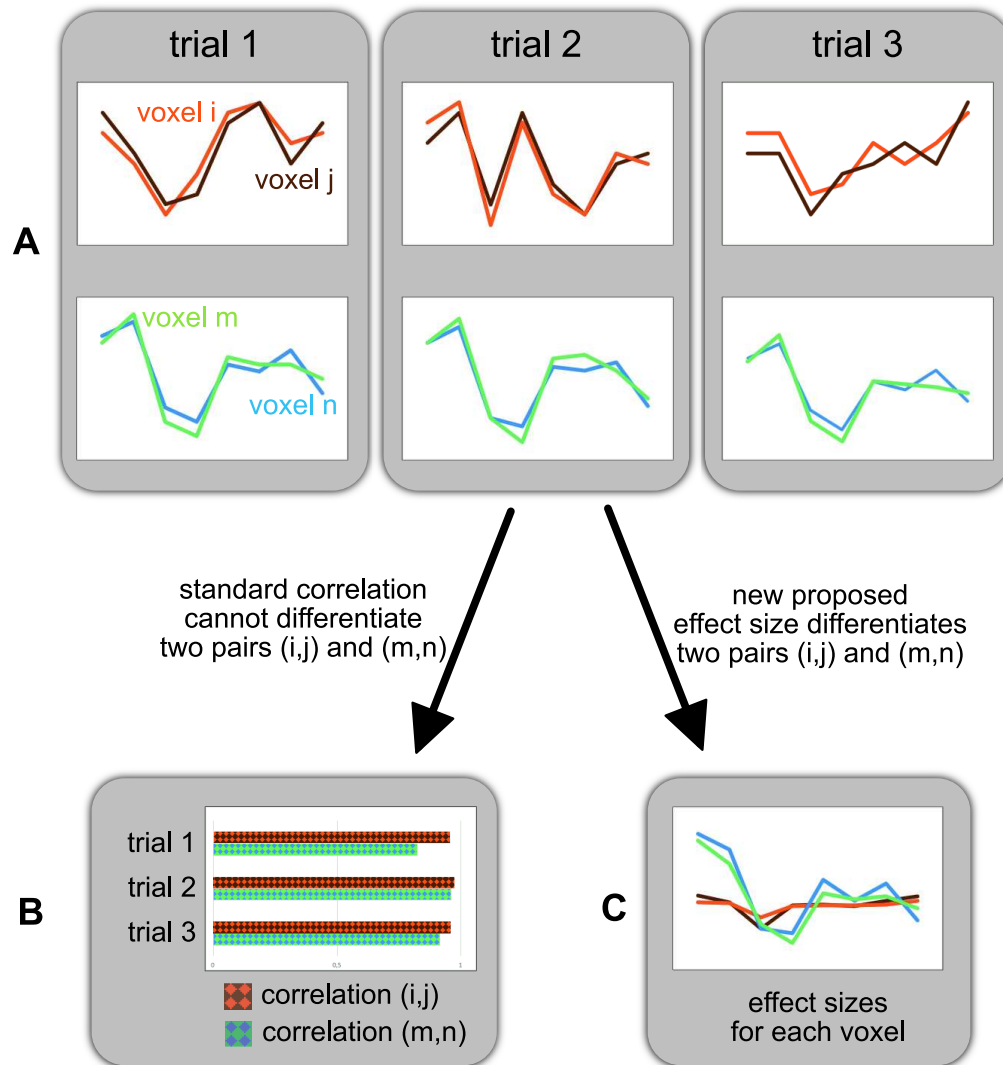


Figure 1. Illustration of a potential problem in correlation-based statistics. (A) Hypothetical time courses of two pairs of voxels (i, j) and (m, n) in three experimental trials of the same condition are shown here. It is clearly visible that the voxel pair (i, j) has a low inter-trial consistency, while the pair (m, n) has a high one. (B) In standard correlation-based statistics, the correlation between pairs of voxels is computed for each trial. Thus all voxel pairs receive high values of correlation and may thus be interpreted as belonging to the same network. Inter-trial consistency is not taken into account here. However, if the connectivity between voxels is truly induced by the task and the response processing is assumed to be similar, the high correlations of the pair (i, j) might be caused by nuisance, while the pair (m, n) reflects a true effect. Clearly, this is not visible using a standard correlation approach. (C) We propose a new measure of synchronization based on effect sizes, taking into account the inter-trial consistency. Our measure is able to separate between the voxel pairs (i, j) and (m, n); the voxel pair with low inter-trial consistency receives low scores.

interaction (PPI) method [20] and its generalisation [21], where the interaction usually is computed after deconvolution of the fMRI signal into the neural space [22]. Another prominent example is the *beta correlation method* which detects the correlation of parameter estimates from the seed area to the rest of the brain [23]. The parameter estimates themselves are derived from a general linear model (GLM).

The weak point of seed-based methods is their inability to reveal *global* changes of functional reorganisation. Only differences *relative to the seed area* can be depicted so that just a small part of the picture is revealed. Thus, a full exploration would require a multitude of seed-based analyses (i.e. one for each grey matter location) and to combine the resulting maps in a second step. It is easy to see that such a procedure constitutes a daunting multiple comparisons problem, which ultimately renders a whole-brain approach infeasible. A further problem arises from analysing *correlations* in time series, as differences in correlations are in general not very reliable indicators of membership in a network. Indeed, some time series may show very similar correlations and yet belong to very different networks. For a graphical illustration of the issue, see Fig. 1.

The choice of seed areas raises further issues, as only few seed locations can be studied without a proper multiple comparisons correction. Researchers thus are required to carefully select the locations in question. It is common practice to use seed locations of special anatomical interest, or alternatively, to choose seed locations of activation peaks (e.g. as determined by a prior whole-brain GLM analysis). However, the latter procedure assumes that brain regions featuring relevant changes in terms of connectivity are indeed also activation peaks. This assumption has recently been challenged empirically [24].

An alternative way of performing network-based analyses is to use parcellation schemes, reducing the number of network nodes. For instance, *condition-specific networks* [25] reveals changes in the whole-brain connectivity structure that occur as response to a task, depicting the variation of functional connectivity between pairs of regions. Similarly, *network-based statistics* [26] evaluates changes in the network structure and incorporates a solution to the multiple comparisons problem based on a graph-based connected components methodology. Thus, this method principally allows a larger number of smaller regions. Further methods include the adaptation of PPI on parcellations [24], which enables the investigation of global changes. At the other end of the spectrum in terms of involved brain regions stands *dynamic causal modelling* (DCM), which attempts to investigate causal influences within very small networks [27]. DCM's validity was challenged in [28].

Parcellating the brain into regions comes with a number of issues, however. The choice of the parcellation scheme underlies a certain degree of arbitrariness and is rarely motivated by anatomical features such as cyto- or myeloarchitecture [29]. This gives rise to nonlinear properties, where small

deviations in the size of regions can result in large changes in underlying network connectivity [30]. Therefore it is no surprise that the choice of parcellation scheme and thus the number of regions have a substantial impact on the resulting network metrics [31, 32, 33, 34]. Furthermore, region-based approaches assume functional homogeneity within the regions [35]. This is particularly troubling if the regions are large enough so that they can be further subdivided into parts that feature heterogeneous connectivity profiles (e.g. see [36, 37, 38, 39, 40]). Averaging within such heterogeneous regions may effectively hamper the detection of subtle connectivity changes that occur only in a subregion. Furthermore, it should be mentioned that adopting a parcellation scheme also implies that it is not possible to quantify the total number of connections between regions [41]. Therefore, it has been suggested that voxel-level approaches in the context of network analysis are preferable [35].

Several other algorithms target only network hubs rather than entire networks, e.g. [42, 43] and are therefore not comparable to the present approach. Other methods such as multivariate pattern analysis (MVPA) [44] or independent component analysis (ICA) [45] also fall into a different domain, and are therefore not discussed here.

The publications listed above all contributed immensely towards a network-based understanding of brain function. However, they all suffer from some limitations, e.g. they require a presegmentation of the data, or they do not offer a mechanism for statistical inference, or they depend on a particular hemodynamic model. The dependence on a hemodynamic model was found to be problematic in a recent study by Gonzalez-Castillo et al. who tested a range of different hemodynamic response models and found wide-spread activations which had previously evaded detection [46]. They ascribed the sparsity of classical activation maps to high noise levels and overly strict response models.

Therefore, our goal in this paper is to establish a new method for fMRI data analysis that fulfills the following requirements. It should

1. identify task-related changes in network configuration,
2. not require any presegmentations,
3. be free from any specific hemodynamic response model,
4. and incorporate rigorous statistical inference.

To achieve this goal, we characterise functional networks as large-scale, task-related *collective synchronisations* of the BOLD signal measured at voxel-level resolution. At the heart of our method is the concept of spatially localised and task-related edge density motivating us to call this algorithm "TED" (Task-related Edge Density). In short, TED identifies edges in a brain network that differentially respond in unison to a task onset and that occur in dense packs with similar characteristics. We found TED to be a very strong marker for dynamic network formation that easily lends itself to statistical analysis using large scale statistical approaches.

In the following, we describe our proposed algorithm and demonstrate its applicability for dynamic network discovery in task-based fMRI data provided by the Human Connectome Project (HCP) [47].

2. Materials and Methods

Networks consist of nodes which are interconnected by edges. We define nodes as voxels and the weight of an edge between any pair of voxels as *task-related changes in dynamic synchronisation* between their respective time series.

Our algorithm supports experiments presented in a block design with non-overlapping trials of sufficient length to allow for a connectivity analysis. It computes a change in synchronization between two experimental conditions and it requires several repetitions of trials of both conditions to permit a valid statistical inference.

The algorithm proceeds in six steps. First, the data are preprocessed using a standard preprocessing pipeline which must include a correction for baseline drifts. Second, we define a measure z of task-related differential synchronisation for each edge in the network. In a third step, the z -values are normalised. Fourth, a measure called “edge density” (D_e) is computed for each edge, after which edge densities are subjected to a statistical inference procedure to assess which edges are significantly affected by the experimental task. Finally, we propose methods for visualising the results. In the following, we will describe the six processing steps of TED in more detail.

Step 1. Preprocessing.

The TED algorithm assumes that the fMRI data have been preprocessed using some standard preprocessing pipeline. This should generally include corrections for motion, slicetiming, and EPI-related distortions as well as a removal of baseline drifts. In the case of multi-subject studies, a geometric alignment with the MNI anatomical template is needed. Physiological noise removal should be included into the preprocessing chain if there is reason to assume that it differentially affects the two task conditions.

Step 2. Obtaining a measure z of task-related differential synchronisation.

Let A and B denote two experimental conditions such as left hand versus right hand fingertapping presented in a block design. In our experiments, the trial duration was 12 seconds.

For simplicity, we assume that all trials have the same duration T , and there are K number of trials per condition. For condition A , let $v_i^A(k, t)$ denote the time course of voxel i of trial k at time t . We now define a measure that quantifies the amount of task-related change in connectivity between any two voxels i and j . A seemingly straightforward way, such as correlation-based statistic (CBS), would be to simply compare the linear correlations of their respective time series during the execution of different tasks. However, CBS can be problematic because it assumes that similarities in correlations are sufficient for assuming membership in common

network, which is problematic (see Fig. 1). In the following, we therefore propose a different measure which we call *differential synchronisation* z .

We first compute the average μ and standard deviations σ across all trials as follows.

$$\mu_i^A(t) = \frac{1}{K} \sum_{k=1}^K v_i^A(k, t) \quad (1)$$

$$\sigma_i^A(t) = \sqrt{\frac{1}{K-1} \sum_{k=1}^K (v_i^A(k, t) - \mu_i^A(t))^2} \quad (2)$$

For each voxel i we thus obtain an effect size at time point t :

$$s_i^A(t) = \frac{\mu_i^A(t)}{\sigma_i^A(t)}$$

The synchronisation $\theta_{i,j}^A$ between voxels i and j in condition A is then defined as the z -transformed linear correlation between $s_i(t)$ and $s_j(t)$. More precisely, we have

$$\theta_{i,j}^A = \begin{cases} \frac{1}{2} \log \frac{1+r_{i,j}^A}{1-r_{i,j}^A} & \text{for } r_{i,j}^A > 0 \\ 0 & \text{otherwise} \end{cases}$$

with $r_{i,j}^A$ denoting the Pearson correlation coefficient between s_i^A and s_j^A .

The synchronisations for experimental condition B are computed analogously.

We now have two $n \times n$ matrices Θ_A and Θ_B each recording the task-related synchronisation in conditions A, B for all pairs of voxels $i, j = 1, \dots, n$ where n is the number of voxels. Based on these two matrices, we introduce a measure of differential synchronisation z defined as an elementwise subtraction:

$$z_{i,j} = \theta_{i,j}^A - \theta_{i,j}^B \quad (3)$$

Note that large positive values of z indicate a higher synchronisation in condition A compared to condition B , see Fig. 2 for an illustration.

Negative correlations are excluded in the definition of θ to avoid misinterpretations. Specifically, consider a case where two voxels are not correlated at all in experimental condition A , while showing a strong negative correlation in condition B . If the synchronisation θ were allowed to take negative values, then this would entail $z = \theta^A - \theta^B \gg 0$ which might be mistaken for a task-positive involvement of the connection between these two voxels in condition A , even though the correlation is in fact absent.

Also note that z is designed to enforce task-related synchrony across trials within the same experimental condition so that the problem illustrated in Fig. 1 does not arise.




Figure 2. Illustration of differential synchronisation $\theta_{i,j}$. The figure shows mean $\mu(t)$ and standard errors $\sigma(t)$ across 100 trials in a pair of voxels i and j . The left pane shows experimental condition A, the right pane condition B. The two voxels of condition A appear to be stronger synchronised than those of condition B. This is reflected by a high value of synchronisation in condition A which is $\theta_{i,j}^A = 1.453$ while for B it is only $\theta_{i,j}^B = 0.108$. The time courses are taken from two voxels of the experimental data described in this article.

Step 3: Normalisation of z -values.

The output of the previous step is an $n \times n$ symmetric matrix of z -values. With $n \approx 54,000$, this matrix has more than one billion entries. Theoretically, we might want to apply a statistical test to the elements of this matrix in order to find significant task-related differences in synchronisation between two voxels. However, such an approach would be problematic because the z -values must be expected to be heavily influenced by non-neuronal confounds such as cardiac and respiratory effects or subject motion [48, 49]. Furthermore, it is well known that neuromodulatory effects (state of arousal, attentiveness, etc.) have a major impact on the fMRI signal [50]. It is extremely difficult to disentangle “true” neuronal effects from the confounding effects listed above. Therefore - instead of trying to solve this challenging problem - we propose to evade it using normalisation. Specifically, we choose the following approach.

We begin by computing the probability density distribution of the z -values. Theoretically - if no effect were present in the data - this distribution should be a Gaussian normal. However, due to confounds, we cannot expect this to be the case. Therefore, we apply a histogram matching procedure so that this distribution does indeed follow a Gaussian normal with mean zero and standard deviation one. We call this procedure z -normalization and it can be easily achieved by any standard histogram matching algorithm [51].

At this point, it would not make sense to apply an elementwise statistical test to the normalized z -values to check for significant differences from zero. The reason is that after normalization the distribution of the z -values is exactly the same as that of the theoretical null distribution. In other words, normalization effectively eradicates elementwise effects that might have been present in the z -values - regardless

of whether or not they were neuronal or non-neuronal in origin. Clearly, this step makes our approach very conservative. But since normalization is a monotone transformation it preserves the ranks of the z -values, and it also preserves spatial information so that small neighbourhoods containing mostly high ranking z -values will also have high ranking normalized z -values. In the following, we will solely exploit information of this type.

Step 4: Edge densities.

We now propose a new network metric that draws on spatial adjacency as the key source of information. We call this feature “edge density” (D_e). The edge density is computed for all edges in the graph that surpass an initial user-defined threshold z_t . The value of $D_e(i, j)$ for an edge connecting two spatially separate voxels i, j indicates to what degree the two neighborhoods of the voxels i and j are connected with each other. A high edge density indicates that many edges connect the two neighbourhoods, while a low edge density indicates that only few links are present, for an illustration see Fig. 3.

Quantitatively, the edge density is defined as follows: first, the total number of possible edges between the neighborhoods of the voxels i and j is computed (omitting local connections, i.e. the start and ending point of an edge must be in different neighborhoods). Next, the number of supra-threshold edges between the neighborhoods is determined, only taking into account edges whose normalised z -value exceed the threshold z_t . The fraction between this number of suprathreshold edges and the total possible number of edges then defines the edge density D_e . Thus, if all neighbouring edges have supra-threshold z values, D_e will be one, and if most edges fall below the threshold, D_e scores will approach zero. To summarize, D_e indicates to what degree the local neighbourhoods of pairs of voxels show a similar change of connectivity across




Figure 3. Schematic illustration of edge density (D_e), showing a case with high D_e and one with low D_e . In this figure, voxels are depicted as small spheres. The value of D_e of the (thick red) edge connecting the two blue voxels in the centre is computed as follows: First, we consider the 26-adjacent neighbours of its endpoints (shown here as grey and green spheres). Theoretically, the highest possible number of edges connecting any two endpoint voxels *across* the 26-neighbourhoods is $27 \times 27 = 729$. We define the edge density D_e as the number of edges whose z -values are above a user-defined threshold z_t divided by the total number of possible edges (i.e. 729). In the above examples, lines connecting nodes (voxels) indicate supra-threshold edges. In example (A), 11 out of 729 possible edges are above threshold, thus $D_e = \frac{11}{729} \approx 0.015$. In example (B), only 5 out of 729 possible edges are above threshold, thus, $D_e = \frac{5}{729} \approx 0.007$.

the experimental conditions. As the edge density is only computed for edges whose z is larger than z_t , all other edges are defined to have an D_e of zero. In our experiments, we considered the top 1 percent of all edges so that z_t was set to 2.33. Note that the computation of D_e requires a specification of adjacency. In our experiments, we used 26-neighbourhoods, but 18- or 6-adjacencies may also be considered.

Also note that we exclude short edges from further analysis. A short edge is an edge whose endpoints i, j have a Euclidean distance of less than 15mm. The reason for doing this is that the two neighbourhoods should be non-overlapping, and because of spatial smoothness we additionally increased the minimum distance so that the borders of the two neighbourhoods are at least three voxels apart.

Step 5: Statistical inference.

In order to correct for multiple testing we employ a procedure controlling the false discovery rate (Fdr) [52, 53, 54].

The original Fdr algorithm proposed in [52] requires that data points are independent and that the null distribution is uniform. Both requirements may not be fulfilled in our case. Therefore, we use a different formulation of Fdr that is well suited for large-scale statistics involving a large number of data points with complex dependencies among them [53, 54].

The basis of this Fdr approach is the assumption of a two component mixture model for the D_e -scores based on a null and non-null component with a cumulative distribution function (cdf) $F_0(D_e)$ and $F_1(D_e)$ respectively. Additionally, we define an a priori probability for being null by π_0 . Then, the observed joint density $F_{joint}(D_e)$ is given by

$$F_{joint}(D_e) = \pi_0 \times F_0(D_e) + (1 - \pi_0)F_1(D_e) \quad (4)$$

The Fdr is defined as the probability of being null, or a false discovery, given a D_e -score as big or bigger than the observed one. This translates into

$$Fdr(D_e) = \pi_0(1 - F_0(D_e))/(1 - F_z(D_e)). \quad (5)$$

For simplicity we assume $\pi_0 = 1$, which is the most conservative choice. The cdf F_{joint} is estimated from histogram counts of the edge densities using cumulative summation. As we do not have evidence for a theoretical null distribution we rely on an empirical permutation null estimate, where the null cdf F_0 is estimated by using random permutation of task labels [55]. More precisely, we randomly construct a binary permutation vector ρ of size K where each entry $k \in 1, \dots, K$ indicates whether or not the task label should be swapped in trial k . We used Bernoulli random trials with probability $p = 0.5$ for this purpose. This yields

$$\rho(k) = \begin{cases} 0 & : \text{k-th trial original} \\ 1 & : \text{k-th trial swapped} \end{cases}$$

The permuted group mean amplitudes for the experimental conditions A and B for voxel i then are defined as

$$\mu_i^A(t) = \frac{1}{K} \sum_{k=1}^K (1 - \rho(k))v_i^A(k, t) + \rho(k)v_i^B(k, t)$$

$$\mu_i^B(t) = \frac{1}{K} \sum_{k=1}^K \rho(k)v_i^A(k, t) + (1 - \rho(k))v_i^B(k, t)$$

with the standard deviations for experimental condition A and B

$$\sigma_i^A(t) = \sqrt{\frac{1}{K} \sum_{k=1}^K [(1 - \rho(k))v_i^A(k, t) + \rho(k)v_i^B(k, t) - \mu_i^A(t)]^2}$$

$$\sigma_i^B(t) = \sqrt{\frac{1}{K} \sum_{k=1}^K [\rho(k)v_i^A(k, t) + (1 - \rho(k))v_i^B(k, t) - \mu_i^B(t)]^2}$$

In other words, for the special case of the permutation vector $\rho(k) = 0, k = 1 \dots K$, the above definitions simplify to the original definitions (1) and (2).

In order to ensure that spatial smoothness is preserved we apply the same permutation vector ρ to all voxels in the brain mask. Therefore, differences between F_0 and F_z cannot be attributed to the spatial correlations that are generally inherent in fMRI data. As described in more detail below, a relatively small number of permutations may suffice to converge to a stable estimation of F_0 , and hence of the false discovery rate.

Step 6: Visualisation of results.

The previous processing step yields a set of edges that indicate significant changes in task-related connectivity. Since the number of such edges can be very large, visualisation of results may become difficult. Here we propose two different methods.

The first method is to project edges onto a “hubness map”. A voxel in the hubness map records the number of edges for which this voxel serves as an endpoint. Voxels in which many edges accumulate may be viewed as hubs in a task-specific network, and the number of edges meeting in a voxel is a measure of the voxel’s hubness, see Fig. 5. Note however, that this measure of hubness should not be confused with activation strength as can be seen from figure 2. Here differential synchronisation goes along with a decrease in BOLD activation rather than an increase.

The second method is to display edges as lines in a 3D rendering such that each line represents an edge that survived significance thresholding. Such renderings can become quite cluttered and therefore edges that are close to each other are bundled together to produce a clearer picture [56]. We use the software package “braingl” for this purpose [57], see Fig. 6.

Experimental Data

We applied TED to task-based fMRI data provided by the Human Connectome Project (HCP), WU-Minn Consortium [47, 58]. We focused on the motor/fingertapping task using minimally preprocessed fMRI data of 100 participants. The pre-processing protocol is described in [59]. The experiment was acquired in two separate runs with one run using left-right phase-encoding, and the other run using right-left phase-encoding. While in the scanner, participants were cued visually to tap their left or right fingers, squeeze their left or right toes, or move their tongue. Each block lasted 12 seconds (10 movements), and was preceded by a 3 second cue. Since the repetition time was 720 milliseconds, there were 16.666 volumes (time steps) per trial of which we used the initial 16.

In each of the two runs, there were 13 blocks, with 2 of tongue movements, 4 of hand movements (2 right and 2 left), 4 foot movements and three 15 s fixation blocks per run, see [58]. Here, we only used the second of the two finger-tapping blocks so that we have 100 trials for each condition *A* and *B* (right hand tapping and left hand tapping) in each phase-encoding run.

In order to reduce the number of voxels and hence the computational load, we downsampled the data to isotropic voxels of size $(3.0\text{mm})^3$ from the original resolution of $(2.0\text{mm})^3$. We also corrected for baseline drifts using a highpass filter with a cutoff frequency of 1/90 Hz. To reduce the effect of anatomical variability across subjects, we applied spatial smoothing with a Gaussian smoothing kernel of FWHM=5mm. Spatial smoothing is however not an integral part of the algorithm and should be omitted whenever possible, particularly for single-subject analysis [60]. We manually defined a region of interest (ROI) containing about 54,000 voxels covering the entire brain including grey and white matter, subcortical structures, CSF and the cerebellum. Furthermore, we normalised the voxel-wise time series of each trial, so that their time series have mean zero and standard deviation one.

We performed the initial three steps of TED separately for each of the two phase-encoding runs resulting in two matrices of normalised z -values. These two matrices were then combined via a conjunction analysis by taking the element-wise minimum of the z -values. For the subsequent edge density computation we used a threshold of $z_t = 2.33$.

We applied TED as described above using 1000 random permutations to estimate the null distribution. The TED algorithm is implemented in C/C++ and makes use of parallel computation. With 54,000 voxels TED requires about 13 GByte of main memory, and one permutation takes about three minutes of computation time on a Linux PC using 12 parallel cores. For the computation of all 1000 permutations we made use of the Max-Planck-Society’s high-performance computing center in Garching, Germany.

3. Results

We first contrasted global connectivity related changes for right hand minus left hand tapping. Figure 4 shows the distri-

butions F_z and F_0 that were used to estimate statistical significance. At a false discovery rate of $Fdr < 0.05$ the edge density cutoff was found to be $D_e = 0.1621$, i.e. edges with $D_e > 0.1621$ can be assumed to indicate a significantly stronger synchronisation in left hand versus right hand tapping. Corresponding plots for the reverse contrast can also be found in figure 4. Note that far fewer than 1000 permutations would have sufficed to reach a similar result so that the entire computation could actually have been performed on a standard PC.

Figure 5 shows a resulting hubness map produced as described above in step 6. Voxels that are colour-coded are endpoints in an edge significantly affected by the task. Voxels in which many edges accumulate may be viewed as “hubs” in a task-specific network, and the number of edges meeting in a voxel is a measure of the voxel’s “hubness”. The upper panel indicates the hubness for edges in the contrast right hand minus left hand. Tapping with the right hand as opposed to the left increased the global connectivity in supplementary motor areas, right and left motor cortex, somatosensory areas, the frontal eye fields, regions in the parietal cortex and the visual cortex. On the other hand, as shown in the lower panel, left hand tapping minus right hand tapping seemed to increase the global connectivity within the bilateral motor network, the default mode network, bilateral putamen, bilateral V5, insular cortex and regions in the cerebellum.

In the same figure it can be seen that several regions appear involved in both contrasts. The reason is that a single brain region may participate in two different networks that are differentially activated by the tasks. Note that such a case cannot be detected in a classical GLM-based analysis. As an example, we investigated the right primary motor cortex which is involved in both contrasts, see Figs. 6,S3. Here we only investigated edges with one endpoint in a preselected region of interest, which roughly corresponds to the right hand area. The maps reveal a striking difference in synchronisation of the right primary motor cortex to the rest of the brain: the right hand tapping condition involves stronger synchronisation between the right motor cortex and regions in the visual and parietal cortex. In the left hand condition, the same area shows higher synchrony with bilateral areas in the cerebellum, V5, the putamen and insular cortex and furthermore the medial prefrontal cortex.

Comparison with a GLM-based analysis.

For comparison, we performed a standard analysis using the GLM approach as implemented in *Lipsia* [62]. The pre-processing of the data was performed as described above. We computed activation maps for each of the two phase-encoding runs separately using the general linear model. These maps contain uncorrected z -values representing the contrast between left hand minus right hand fingertapping. As in the TED approach, we performed a conjunction analysis on the two maps, where the voxel-wise minimum value of both Z -maps was used for the case that both were positive, the maximum

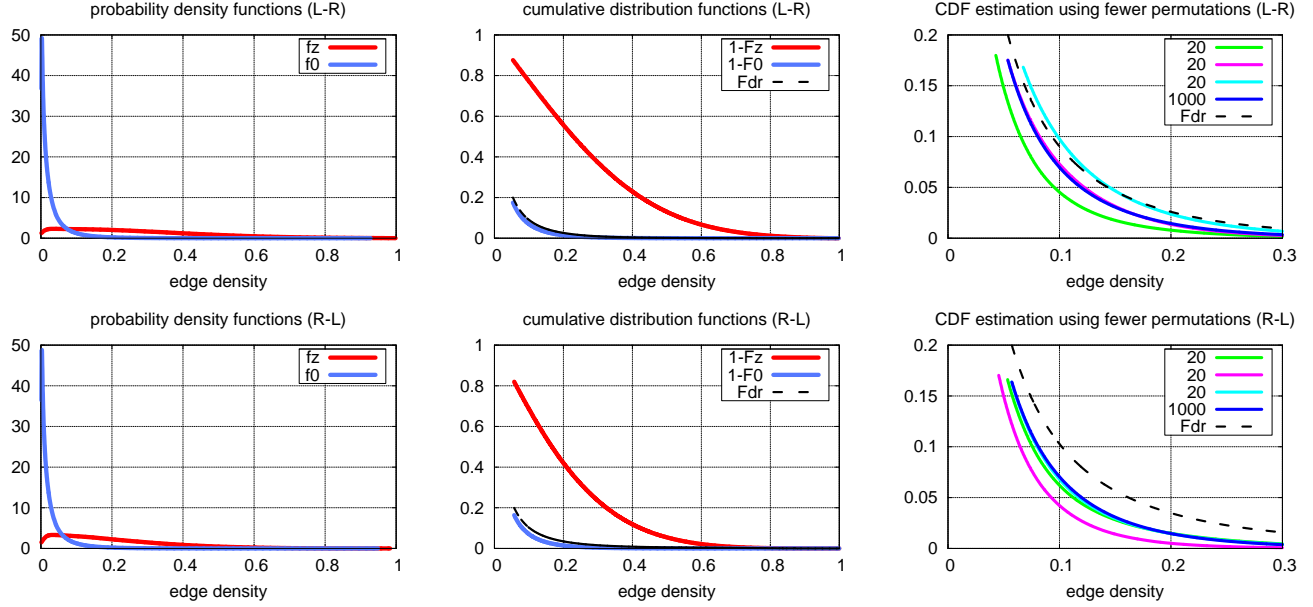


Figure 4. Probabilities of the edge density estimated from fMRI fingertapping data. The top rows contrasts left hand minus right hand tapping, the bottom row is the reverse contrast (right hand minus left hand). The left plots show estimations of the probability density functions f_0 (permutation-derived) and f_z (no permutation) using Gaussian kernels with kernel bandwidth defined via Silverman’s rule [61]. The plots in the middle show the corresponding cumulative distribution functions (CDF) F_0 and F_z . Note that the differences between the null and the “real” distributions are massive. The CDFs are used to estimate the false discovery rates (Fdr) which are indicated here by dashed black lines. For better visualization, the same Fdr curves are also shown in the plots on the right. The Fdr is used to determine a significance threshold. For the left minus right contrast, the cutoff was found to be $D_e > 0.1435$, i.e. for edges with $D_e > 0.1435$ the false discovery rate falls below 0.05. For the reverse contrast, the cutoff was $D_e > 0.1621$. The estimation of F_0 is based on 1000 random permutations. The right plots show estimations of F_0 using only 20 different permutations. Note that the plots are zoomed for better visualisation. Here we show results of three random selections of such shorter permutations vectors. Note that they are very similar to the estimation based on 1000 permutation leading to very similar estimations of Fdr. In other words, far fewer than 1000 permutations would have sufficed to reach a similar result.




Figure 5. Task-dependent dynamic reconfiguration of whole-brain networks. We depict the reconfiguration using hubness maps on basis of fMRI data fingertapping data of the Human Connectome Project. The hubness maps indicate the number of network edges that feature a significant change between the two experimental conditions. The top row (A) contrasts right hand minus left hand tapping, the bottom row (B) shows the reverse contrast. The colours encode the number of edges with $Fdr < 0.05$ having one of their endpoints in the respective colour-coded voxel and ranges from 1 to 1000. This number can be interpreted as a measure of “hubness”. Thus, red values in the above figure indicate hubs where many edges accumulate in a voxel. See also supplementary Figs. S1,S2.




Figure 6. Task-related reconfiguration of the right-hemispheric primary motor cortex. In a TED analysis, a single region may appear in opposite contrasts because it may participate in different task-dependent networks. In this figure we show the participation of the right primary cortex (hand knob area, marked with a green box) in different networks, depending on the experimental condition. The exact shape of the region of interest is shown in supplementary figure S3. In the upper panel we display voxels involving all edges featuring a stronger synchronisation with the right primary motor cortex for right hand fingertapping (as compared to left hand tapping). In the lower panel we show the voxels where the synchronisation with the same area is higher for left hand fingertapping (as opposed to right hand tapping). The maps reveal a striking difference in synchronisation of the right primary motor cortex to the rest of the brain: the right hand tapping condition involves stronger synchronisation between the right motor cortex and regions in the visual and parietal cortex. On the other hand, in the left hand condition the same area synchronizes more with bilateral areas in the cerebellum, V5, the putamen and insular cortex and furthermore the medial prefrontal cortex. Below, the same data are shown using a 3D rendering using the software package “braingl” [57, 56]. Here, the synchronisation network of the right primary motor cortex (the red sphere) in the right hand minus left hand contrast is shown in green, the reverse contrast is shown in yellow.




Figure 7. Classical univariate GLM-based analysis. For comparison to standard analysis methods, we used a univariate activation based GLM technique (see methods). We thresholded the activation map very liberally at $|z| > 2.33$ on the voxel level without correcting for multiple comparisons. See also supplementary Fig. S4.

value in case both were negative, and zero for diverging signs of the z -values. We then thresholded the resulting conjunction map such that voxels with $|z| > 2.33$ remained. No multiple comparisons correction was applied. The resulting map contrasting left hand minus right hand fingertapping (see Fig. 7) shows the voxel-wise differences in activation strength. The activations include bilateral motor areas, pre-SMA, as well as the cerebellum.

4. Discussion

In this study, we introduced a new algorithm called “TED” which was designed to identify task-related reconfigurations in brain networks without requiring presegmentations and without being dependent on some specific hemodynamic response function. Since network-based analysis methods face tremendous multiple comparison problems, statistical inference was a key concern in this context.

At the heart of TED is the concept of “edge density”. A voxel is deemed to be involved in the task if it has a partner voxel at some spatially distinct location with a similar synchronicity trajectory, so that these two voxels form an edge with a supra-threshold differential synchronisation. At the same time, such an edge is required to appear in a dense pack of neighbouring edges with similar characteristics. This approach freed us from the need for any explicit hemodynamic modelling, and it also allowed us to make inferences at the spatial resolution of small neighbourhoods around individual voxels without requiring a presegmentation.

TED’s validity is tested via a comparison against a null model derived from permutation testing. We found that the null and the real distributions of edge densities differ massively so that statistical significance could be easily established. The amount of this difference is especially surprising since it is based on normalized z values so that the size of this effect can only be ascribed to spatial adjacency, not to

the magnitudes of the z values themselves. We normalized the z values to avoid potential confounds from physiological noise. Spatial structure was preserved in our random permutations so that our results cannot simply be due to the spatial smoothness that is generally inherent in fMRI data.

The results obtained with TED suggest a dominant role of local inter-connected neighbourhoods forming transient task-related networks with other local neighbourhoods. In this regard, our concept of edge density is somewhat related to the concepts of a clustering coefficient and small-worldness [63, 2]. However, edge density differs from the clustering coefficient in that it measures connectivity between two local neighbourhoods at spatially separate areas. In the literature, there exist several approaches comparing the differences between local and global functional connectivity profiles [64, 65]. It has been found that different brain regions exhibit a varying balance between such local and global connectivity, which may further depend on the current experimental state (i.e. task). On the other hand, the results of both Sepulcre et al. [64] and Tomasi et al. [65] indicate a strong overlap between regions which feature both increased local and global functional connectivity. Our methodology offers a potentially interesting perspective on this, as according to our interpretation global changes in connectivity may be accompanied by local interconnections, reflecting the dynamics of local sub-networks that form transient long-range connections.

The hubness maps produced by TED are generally consistent with the standard univariate GLM analysis and also with a meta-analysis on fingertapping [66]. However, TED found several sites that significantly changed their task-related TED trajectories while not reaching significant net-BOLD modulation when analysed using the traditional GLM approach. This agrees with earlier findings by Gerchen et al. [24]. Also, some regions with negative task involvement showed positive TED changes, and vice versa. This shows that brain areas that appear non-significant in a GLM activation map may

nevertheless form relevant hubs in a task-positive network, and that the sign of task-related activity can be independent of the sign of task-related (de-)synchronisation. TED-based fMRI analysis thus constitutes a novel analysis approach that complements traditional GLM based analyses. In particular, it can reveal task-related regional involvement that evades detection using traditional approaches.

In this respect, our work is in line with an earlier study by Gonzalez-Castillo et al. who also found large-scale time locked activity that went undetected in the classical fMRI analysis [46]. The authors ascribed the sparsity of traditional activation maps to high noise levels and overly strict predictive response models. Our present study goes beyond this earlier work in that we present data-driven criteria that can be subjected to rigorous statistical significance testing. Also, the amount of scan time required to achieve this result was considerably smaller than in this earlier study. Here we only needed 4×100 trials where each trial was 12 seconds long. Adding a hypothetical intertrial interval of perhaps 18 seconds we arrive at a scan time of roughly 200-minutes.

In the present study we applied TED to group level data so that spatial accuracy was limited. This was even further reduced because we had to downsample the data to $(3\text{mm})^3$ resolution to ease the computational burden, see [60] for a discussion on problems relating to spatial inaccuracies. These limitations are not implicit in the TED algorithm, so that it is quite possible to apply TED to single subject data at very high spatial resolutions provided sufficient computational resources exist and a sufficient number of experimental trials are acquired to yield sufficient statistics. Since TED specifically targets local neighbourhoods, we expect that the results will benefit from more precise spatial information, e.g. provided by ultrahighfield MRI scanners.

Note that our present implementation of TED assumes that inter-trial variance is the only source of variation in the data. However, the input data may be structured such that several sources of variance need to be addressed. For instance, both within-subject and between-subject subject variance may be relevant. A potential way of handling such cases may be to perform the first three steps of TED for each subject individually. This results in a separate z-matrix per subject each incorporating within-subject variance. These matrices can then be subjected to a onesample t-test incorporating between-subject variance. This results in a single combined matrix which can then be used as input into edge density computations. Statistical inference should then again be based on permutation testing over task labels applied to edge density values. We think that this approach would combine within-subject and between-subject variation in a logical way. However, developing a statistical framework that fully implements statistical inference with multiple sources of variance is beyond the scope of this paper, and will be the object of future work.

The interpretation of TED-results may not always be straightforward. As can be seen from Fig. 2, synchronicity trajec-

tories can take any shape, and may even show a downward slope following stimulus onset. Therefore, strong hubs in TED hubness maps do not necessarily indicate strong activation. Also, as was the case in the Gonzalez-Castillo study [46] - we now see many more brain regions that appear to be involved in the task. The exact role of many of these regions and their mutual interactions are difficult to assess. Furthermore, a single brain area found by TED may be involved in different task-dependent networks and hence appear task-positive in reversed contrasts. We presented the right-hemispheric motor area as an example and show that this brain region participates in different networks depending on the task (see Fig. 6). Note that a classical GLM analysis is not able to depict such a scenario, even though such a reconfiguration is a highly plausible rendition of human brain function. For example, in the case of fingertapping, we would expect to find effects due to handedness so that the contrast “left hand minus right hand” should not simply be the same as “right hand minus left hand” with the sign reversed. And indeed, the TED hubness maps of the first contrast show a remarkably different pattern from that of the reversed contrast - an effect that may well be ascribed to handedness [67]. In sum, the fact that we now have to deal with entire networks rather than univariate regions adds another level of complexity to data interpretation and visualisation. On the other hand, this complexity likely reflects much more closely the true complexity of human brain function.

Acknowledgments

Data were provided by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

References

- [1] F.A.C. Azevedo, L.R.B. Carvalho, L.T. Grinberg, J.M. Farfel, R.E.L. Ferretti, R.E.P. Leite, W.J. Filho, R. Lent, and S. Herculano-Houzel. Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain. *The Journal of Comparative Neurology*, 513(5):532–541, 2009.
- [2] D.J. Watts and S.H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393(6684):440–442, 1998.
- [3] G. Deco and V.K. Jirsa. Ongoing Cortical Activity at Rest: Criticality, Multistability, and Ghost Attractors. *The Journal of neuroscience*, 32(10):3366–3375, 2012.
- [4] G. Deco, M. Senden, and V. Jirsa. How anatomy shapes dynamics: a semi-analytical study of the brain at rest by a simple spin model. *Frontiers in computational neuroscience*, 6:68, 2012.

- [5] C.J. Honey, R. Kötter, M. Breakspear, and O. Sporns. Network structure of cerebral cortex shapes functional connectivity on multiple time scales. *PNAS*, 104(24):10240–10245, 2007.
- [6] R.B. Buxton, K. Uludag, D.J. Dubowitz, and T.T. Liu. Modeling the hemodynamic response to brain activation. *Neuroimage*, 23:S220–S233, 2004.
- [7] K.J. Friston, A.P. Holmes, J.B. Poline, P.J. Grasby, S.C. Williams, R.S. Frackowiak, and R. Turner. Analysis of fMRI time-series revisited. *Neuroimage*, 2(1):45–3, 1995.
- [8] G. Lohmann, J. Stelzer, J. Neumann, N. Ay, and R. Turner. “More Is Different” in Functional Magnetic Resonance Imaging: A Review of Recent Data Analysis Techniques. *Brain Connectivity*, 3(3):223–239, 2013.
- [9] B. Horwitz. The elusive concept of brain network: Comment on “Understanding brain networks and brain organization” by Luiz Pessoa. *Phys Life Rev*, (in press, 2014). doi:10.1016/j.plprev.2014.06.019.
- [10] O. Sporns. Contributions and challenges for network models in cognitive neuroscience. *Nature Neuroscience*, 17:652–660, 2014.
- [11] L. Pessoa. Understanding brain networks and brain organization. *Phys Life Rev*, (in press, 2014). doi:10.1016/j.plprev.2014.03.005.
- [12] K.J. Worsley, J. Cao, T. Paus, M. Petrides, and A.C. Evans. Applications of random field theory to functional connectivity. *Human Brain Mapping*, 6:364–367, 1998.
- [13] F. Carbonell, K.J. Worsley, and N.J. Trujillo-Barreto. On the Fisher’s Z transformation of correlation random fields. *Statistics and Probability Letters*, 79:780–788, 2009.
- [14] K.J. Friston. Functional and Effective Connectivity: A Review. *Brain Connectivity*, 1(1):13–36, 2011.
- [15] S.M. Smith. The future of FMRI connectivity. *Neuroimage*, 62(2):1257–1266, 2012.
- [16] R.M. Hutchison, T. Womelsdorf, E.A. Allen, P.A. Bandettini, V.D. Calhoun, M. Maurizio Corbetta, S.D. Penna, J.H. Duyn, G.H. Glover, J. Gonzalez-Castillo, D.A. Handwerker, S. Keilholz, V. Kiviniemi, D.A. Leopold, F. de Pasquale, O. Sporns, M. Walter, and C. Chang. Dynamic functional connectivity: Promise, issues, and interpretations. *Neuroimage*, 80:360–378, 2013.
- [17] A. Fornito, A. Zalesky, and M. Breakspear. Graph analysis of the human connectome: Promise, progress, and pitfalls. *Neuroimage*, 80:426–444, 2013.
- [18] G. Lohmann, J. Stelzer, J. Neumann, N. Ay, and R. Turner. “More is different” in fMRI: a review of recent data analysis techniques. *Brain Connectivity*, 3(3):223–239, 2013.
- [19] B. Biswal, F.Z. Yetkin, V.M. Haughton, and J.S. Hyde. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnetic Resonance in Medicine*, 34(4):537–541, 1995.
- [20] K.J. Friston, C. Buechel, G.R. Fink, J. Morris, E. Rolls, and R.J. Dolan. Psychophysiological and modulatory interactions in neuroimaging. *Neuroimage*, 6(3):218–229, 1997.
- [21] D.G. McLaren, M.L. Ries, G. Xu, and S.C. Johnson. A generalized form of context-dependent psychophysiological interactions (gPPI): A comparison to standard approaches. *Neuroimage*, 61(4):1277–1286, 2012.
- [22] D.R. Gitelman, W.D. Penny, J. Ashburner, and K.J. Friston. Modeling regional and psychophysiological interactions in fMRI: the importance of hemodynamic deconvolution. *Neuroimage*, 19(1):200–207, 2003.
- [23] J. Rissman, A. Gazzaley, and M. D’Esposito. Measuring functional connectivity during distinct stages of a cognitive task. *Neuroimage*, 23(2):752–763, 2004.
- [24] M.F. Gerchen, D. Bernal-Casas, and P. Kirsch. Analyzing task-dependent brain network changes by whole-brain psychophysiological interactions: A comparison to conventional analysis. *Human Brain Mapping*, 35(10):5071–5082, 2014.
- [25] S. Dodel, N. Golestani, C. Pallier, V. ElKouby, D. Le Bihan, and J.B. Poline. Condition-dependent functional connectivity: syntax networks in bilinguals. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1457):921–935, 2005.
- [26] A. Zalesky, A. Fornito, and E.T. Bullmore. Network-based statistic: Identifying differences in brain networks. *Neuroimage*, 53(4):1197–1207, 2010.
- [27] K.J. Friston, L. Harrison, and W. Penny. Dynamic causal modeling. *Neuroimage*, 19:1273–1302, 2003.
- [28] G. Lohmann, K. Erfurth, K. Müller, and R. Turner. Critical comments on dynamic causal modelling. *Neuroimage*, 59(3):2322–2329, 2012.
- [29] R. Turner and S. Geyer. Comparing Like with Like: The Power of Knowing Where You Are. *Brain Connectivity*, 4(7):547–557, 2014.
- [30] T. Liu. A few thoughts on brain ROIs. *Brain Imaging and Behavior*, 5(3):189–202, 2011.
- [31] J. Wang, L. Wang, Y. Zang, H. Yang, H. Tang, Q. Gong, Z. Chen, C. Zhu, and Y. He. Parcellation-dependent small-world brain functional networks: A resting-state fMRI study. *Human Brain Mapping*, 30(5):1511–1523, 2009.
- [32] S. Hayasaka and P.J. Laurienti. Comparison of characteristics between region-and voxel-based network analyses in resting-state fMRI data. *Neuroimage*, 50(2):499–508, 2010.

- [33] A. Zalesky, A. Fornito, I.H. Harding, L. Cocchi, M. Yücel, C. Pantelis, and E.T. Bullmore. Whole-brain anatomical networks: Does the choice of nodes matter? *Neuroimage*, 50(3):970–983, 2010.
- [34] M.A. de Reus and M.P. van den Heuvel. The parcellation-based connectome: Limitations and extensions. *Neuroimage*, 80(C):397–404, 2013.
- [35] M.L. Stanley, M.N. Moussa, B.M. Paolini, R.G. Lyday, J.H. Burdette, and P.J. Laurienti. Defining nodes in complex brain networks. *Frontiers in computational neuroscience*, 7(169), 2013.
- [36] E. Solano-Castiella, A. Schäfer, E. Reimer, E. Türke, T. Pröger, G. Lohmann, R. Trampel, and R. Turner. Parcellation of human amygdala in vivo using ultra high field structural MRI. *Neuroimage*, 58(3):741–748, 2011.
- [37] A.E. Cavanna and M.R. Trimble. The precuneus: a review of its functional anatomy and behavioural correlates. *Brain*, 129(Pt 3):564–583, 2006.
- [38] S. Thuret. Why looking at the whole hippocampus is not enough—a critical role for anteroposterior axis, subfield and activation analyses to enhance predictive value of hippocampal changes for Alzheimer’s disease diagnosis. *Frontiers in Cellular Neuroscience*, 8(95), 2014.
- [39] E.C. Cieslik, K. Zilles, S. Caspers, C. Roski, T.S. Kellermann, O. Jakobs, R. Langner, A.R. Laird, P.T. Fox, and S.B. Eickhoff. Is There “One” DLPFC in Cognitive Action Control? Evidence for Heterogeneity From Co-Activation-Based Parcellation. *Cerebral Cortex*, 23(11):2677–2689, 2013.
- [40] Y. Zhang, L. Fan, Y. Zhang, J. Wang, M. Zhu, Y. Zhang, C. Yu, and T. Jiang. Connectivity-Based Parcellation of the Human Posteromedial Cortex. *Cerebral Cortex*, 24(3):719–727, 2014.
- [41] M.N. Moussa, C.D. Vechlekar, J.H. Burdette, M.R. Steen, C.E. Hugenschmidt, and P.J. Laurienti. Changes in cognitive state alter human functional brain networks. *Frontiers in Human Neuroscience*, 5:83, 2011.
- [42] R.L. Buckner, J. Sepulcre, T. Talukdar, F.M. Krienen, H. Liu, T. Hedden, J.R. Andrews-Hanna, R.A. Sperling, and K.A. Johnson. Cortical hubs revealed by intrinsic functional connectivity: Mapping, assessment of stability, and relation to Alzheimers disease. *Journal of Neuroscience*, 29(6):1860–1873, 2009.
- [43] G. Lohmann, D.S. Margulies, A. Horstmann, B. Pleger, D. Lepsien, J. Goldhahn, H. Schloegl, M. Stumvoll, A. Villringer, and R. Turner. Eigenvector centrality mapping for analyzing connectivity patterns in fMRI data of the human brain. *PLoS ONE*, 5(4):e10232, 2010.
- [44] K. Norman, S. Polyn, G. Detre, and J. Haxby. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Sciences*, 10(3), 2006.
- [45] C.F. Beckmann, M. De Luca, J.T. Devlin, and S.M. Smith. Investigations into resting-state connectivity using Independent Component Analysis. *Phil Trans Roy Soc Lond Ser B*, 360(1457):1001–1013, 2005.
- [46] J. Gonzalez-Castillo, Z.S. Saad, D.A. Handwerker, S.J. Inati, N. Brenowitz, and P.A. Bandettini. Whole-brain, time-locked activation with simple tasks revealed using massive averaging and model-free analysis. *PNAS*, 109(14):5487–5492, 2012.
- [47] D.C. Van Essen, S.M. Smith, D.M. Barch, T.E.J. Behrens, E. Yacoub, and K. Ugurbil. The WU-Minn Human Connectome Project: An overview. *Neuroimage*, 80:62–79, 2013.
- [48] J.D. Power, K.A. Barnes, A.Z. Snyder, B.L. Schlaggar, and S.E. Petersen. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *Neuroimage*, 59(3):2142–54, 2012.
- [49] N. Khalili-Mahani, C. Chang, M.J. van Osch, I.M. Veer, M.A. van Buchem, A. Dahan, C.F. Beckmann, van Gerven J.M., and S.A. Rombouts. The impact of “physiological correction” on functional connectivity analysis of pharmacological resting state fMRI. *Neuroimage*, 15:499–510, 2013.
- [50] N. Logothetis. What we can do and what we cannot do with fMRI. *Nature*, 453:869–878, 2008.
- [51] R.C. Gonzalez and R.E. Woods. *Digital Image Processing (3rd Edition)*. Prentice-Hall, Inc., Upper Saddle River, NJ, 2006.
- [52] Y. Benjamini and Y. Hochberg. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J. Roy Stat Soc, Ser B.*, 57:289–300, 1995.
- [53] B. Efron. Size, power and false discovery rates. *The Annals of Statistics*, 35(4):1351–1377, 2007.
- [54] K. Strimmer. A unified approach to false discovery rate estimation. *BMC Bioinformatics*, 9(303), 2008. doi:10.1186/1471-2105-9-303.
- [55] T.E. Nichols and A.P. Holmes. Nonparametric permutation tests for functional neuroimaging: a primer with examples. *Human Brain Mapping*, 15(1):1–25, 2002.
- [56] J. Boettger, A. Schäfer, G. Lohmann, A. Villringer, and D. Margulies. Three-dimensional mean-shift edge bundling for the visualization of functional connectivity in the brain. *Transactions on Visualization and Computer Graphics*, 20(3):471–480, 2014.
- [57] <https://code.google.com/p/brainl/>.
- [58] D.M. Barch, G.C. Burgess, M.P. Harms, S.E. Petersen, B.L. Schlaggar, M. Corbetta, M.F. Glasser, S. Curtiss, S. Dixit, C. Feldt, D. Nolan, E. Bryant, T. Hartley, O. Footer, J.M. Bjork, R. Poldrack, S. Smith, H. Johansen-Berg, A.Z. Snyder, and D.C. Van Essen.

Function in the human connectome: task-fMRI and individual differences in behavior. *Neuroimage*, 80:169–189, 2013.

- [59] M.F. Glasser, S.N. Sotiropoulos, J.A. Wilson, T.S. Coalson, B. Fischl, J.L. Andersson, J. Xu, S. Jbabdi, M. Webster, J. Polimeni, D.C. Van Essen, and M. Jenkinson. The minimal preprocessing pipelines for the Human Connectome Project. *Neuroimage*, 80:105–124, 2013.
- [60] J. Stelzer, G. Lohmann, K. Mueller, T. Buschmann, and R. Turner. Deficient approaches to human neuroimaging. *Frontiers in Human Neuroscience*, 8(462), 2014. doi:10.3389/fnhum.2014.00462.
- [61] B.W. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, 1998. ISBN 0-412-24620-1.
- [62] G. Lohmann, K. Müller, V. Bosch, H. Mentzel, S. Hessler, L. Chen, S. Zysset, and D.Y. von Cramon. LIPSIA - a new software system for the evaluation of functional magnetic resonance images of the human brain. *Computerized Medical Imaging and Graphics*, 25(6):449–457, 2001.
- [63] M.P. van den Heuvel, C.J. Stam, M. Boersma, and H.E. Pol. Small-world and scale-free organization of voxel-based resting-state functional connectivity in the human brain. *Neuroimage*, 43(3):528–539, 2008.
- [64] J. Sepulcre, H. Liu, T. Talukdar, I. Martincorena, B.T. Yeo, and R.L. Buckner. The Organization of Local and Distant Functional Connectivity in the Human Brain. *PLoS Computational Biology*, 6(6):e1000808–15, 2010.
- [65] D. Tomasi and N.D. Volkow. Functional connectivity hubs in the human brain. *Neuroimage*, 57(3):908–917, 2011.
- [66] S.T. Witt, A.R. Laird, and M.E. Meyerand. Functional neuroimaging correlates of finger-tapping task variations: An ALE meta-analysis. *Neuroimage*, 42:343–356, 2008.
- [67] S. Klöppel, T. van Eimeren, V. Glauche, A. Vongerichten, A. Münchau, R.S.J. Frackowiak, C. Büchel, C. Weiller, and Siebner H.R. The effect of handedness on cortical motor activation during simple bilateral movements. *Neuroimage*, 34:274–280, 2007.

5. Supplementary Information




Figure S1. Task-dependent dynamic reconfiguration of whole-brain networks. This map is based on the same data as figure 5 of the main manuscript. It shows the hubness map of the contrast left hand minus right hand fingertapping.




Figure S2. Task-dependent dynamic reconfiguration of whole-brain networks. This map is based on the same data as figure 5 of the main manuscript. It shows the hubness map of the contrast right hand minus left hand fingertapping.




Figure S3. Region of interest in the right motor area. The map shows the region of interest in the right moto area used in figure 6 of the main manuscript.

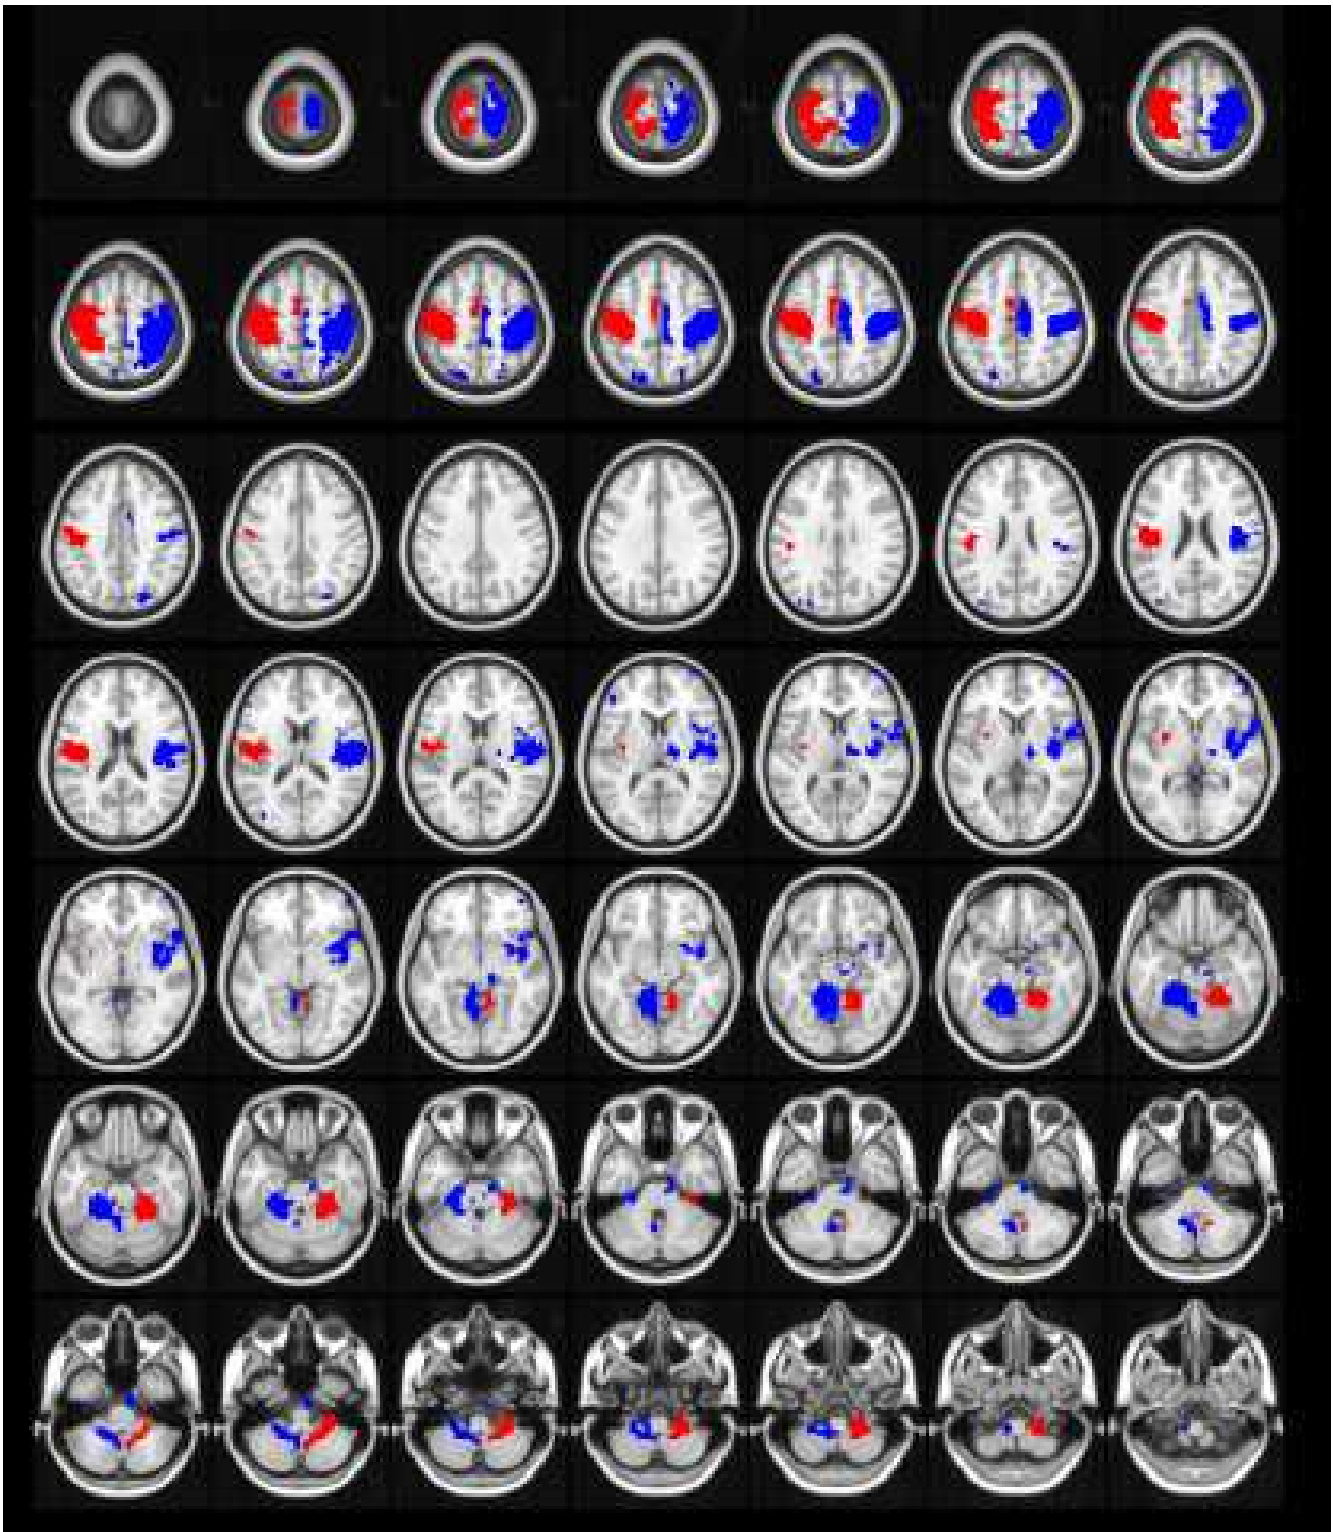


Figure S4. GLM-based activation map. This map is based on the same data as figure 7 of the main manuscript.